On the Eigenvalues of the ADER-WENO Galerkin Predictor

Haran Jackson

Cavendish Laboratory, JJ Thomson Ave, Cambridge, UK, CB3 0HE

arXiv:1611.09153v2 [physics.comp-ph] 19 Dec 2016

Keywords: ADER-WENO, Galerkin, Eigenvalues, Convergence

1. Background

ADER-WENO methods have proved extremely useful in obtaining arbitrarily high-order solutions to problems involving hyperbolic systems of PDEs. For example, it has been demonstrated that for the same computational cost as a Runge-Kutta scheme of a certain order, one can obtain an ADER scheme of one higher order of accuracy (see Balsara et al. [1]). Additionally, Runge-Kutta schemes suffer from the presence of Butcher barriers (see Butcher [3]), limiting the order of temporal accuracy that one can comfortably achieve. There are no such limitations present in ADER-WENO schemes.

The cumbersome analytical derivation of the temporal derivatives of the solution required by the original ADER formulation (see Toro [10]) has been replaced by the use of a cell-wise local Galerkin predictor. The predictor can take either a discontinuous or a continuous form (see Dumbser et al. [4] and Balsara et al. [2], respectively). The Galerkin predictor is a high-order polynomial reconstruction of the data in both space and time, found as the root of a non-linear system.

It has been conjectured that the eigenvalues of certain matrices appearing in these non-linear systems are always zero, leading to desirable system properties for certain classes of PDEs. It is proved here that this is in deed the case for any number of spatial dimensions and any desired order of accuracy, for both the discontinuous and continuous Galerkin variants. This result is independent of the choice of reconstruction basis polynomials.

The Einstein summation convention is to be assumed throughout this paper.

2. The ADER-WENO Method

Take a non-homogeneous, non-conservative (and for simplicity, one-dimensional) hyperbolic system of the form:

$$\frac{\partial \boldsymbol{Q}}{\partial t} + \frac{\partial \boldsymbol{F}(\boldsymbol{Q})}{\partial x} + B(\boldsymbol{Q}) \cdot \frac{\partial \boldsymbol{Q}}{\partial x} = \boldsymbol{S}(\boldsymbol{Q}) \tag{1}$$

where $\boldsymbol{Q}$ is the vector of conserved variables, $\boldsymbol{F}$ is the conservative nonlinear flux, B is the matrix corresponding to the purely non-conservative component of the system, and $\boldsymbol{S}(\boldsymbol{Q})$ is the algebraic source vector.

Take the set of grid points $x_0 < x_1 < \ldots < x_K$ and define $\Delta x_i = x_{i+1} - x_i$. Take the time steps $t_0 < t_1 < \ldots$ while defining $\Delta t_n = t_{n+1} - t_n$. Following the formulations presented in Dumbser et al. [5, 6], Balsara et al. [2], the WENO method and Galerkin method produce at each time step t_n a local polynomial approximation to $\boldsymbol{Q}$ on each space-time cell $[x_i, x_{i+1}] \times [t_n, t_{n+1}]$.

Define the scaled space variable:

$$\chi^i = \frac{1}{\Delta x_i}(x - x_i) \tag{2}$$

Take a basis $\{\psi_0, ..., \psi_N\}$ of P_N and inner product $\langle \cdot, \cdot \rangle$. This basis can either be nodal ($\psi_i(\chi_j) = \delta_{ij}$ where $\{\chi_0, \ldots, \chi_N\}$ are a set of nodal points, such as the Gauss-Legendre abscissae), or modal (such as the Jacobi polynomials).

The WENO method (as used in Dumbser et al. [8]) produces an order-N polynomial reconstruction of the data at time t_n in cell $[x_i, x_{i+1}]$, using $\{\psi_0, \ldots, \psi_N\}$ as a basis. This is denoted:

$$\boldsymbol{w}(x) = \boldsymbol{w}_{\boldsymbol{\gamma}} \psi_\gamma \left(\chi^i(x)\right) \tag{3}$$

This spatial reconstruction at the start of the time step is to be used as initial data in the problem of finding the Galerkin predictor.

Now define the scaled time variable:

$$\tau^n = \frac{1}{\Delta t_n}(t - t_n) \tag{4}$$

Thus, (1) becomes:

$$\frac{\partial \boldsymbol{Q}}{\partial \tau^n} + \frac{\partial \boldsymbol{F}^*(\boldsymbol{Q})}{\partial \chi^i} + B^*(\boldsymbol{Q}) \cdot \frac{\partial \boldsymbol{Q}}{\partial \chi^i} = \boldsymbol{S}^*(\boldsymbol{Q}) \tag{5}$$

where

$$\boldsymbol{F}^* = \frac{\Delta t_n}{\Delta x}\boldsymbol{F} \qquad B^* = \frac{\Delta t_n}{\Delta x} B \qquad \boldsymbol{S}^* = \Delta t_n \boldsymbol{S} \tag{6}$$

Email address: hj305@cam.ac.uk (Haran Jackson)

The non-dimensionalization notation and spacetime cell indexing notation will be dropped for simplicity in what follows. Now define the set of spatio-temporal basis functions:

$$\{\theta_k(\chi,\tau)\} = \left\{\psi_p(\chi)\psi_s(\tau) : 0 \le p, s \le N\right\} \tag{7}$$

Denoting the Galerkin predictor by $\boldsymbol{q}$, take the following set of approximations:

$$\boldsymbol{Q} \approx \boldsymbol{q} = \theta_\beta \boldsymbol{q_\beta} \tag{8a}$$
$$\boldsymbol{F}(\boldsymbol{Q}) \approx \theta_\beta \boldsymbol{F_\beta} \tag{8b}$$
$$B(\boldsymbol{Q}) \cdot \frac{\partial \boldsymbol{Q}}{\partial \chi} \approx \theta_\beta \boldsymbol{B_\beta} \tag{8c}$$
$$\boldsymbol{S}(\boldsymbol{Q}) \approx \theta_\beta \boldsymbol{S_\beta} \tag{8d}$$

for some coefficients $\boldsymbol{q_\beta}, \boldsymbol{F_\beta}, \boldsymbol{B_\beta}, \boldsymbol{S_\beta}$.

If $\{\psi_0, ..., \psi_N\}$ is a nodal basis, the *nodal basis representation* may be used:

$$\boldsymbol{F_\beta} = \boldsymbol{F}\left(\boldsymbol{q_\beta}\right) \tag{9a}$$
$$\boldsymbol{B_\beta} = B\left(\boldsymbol{q_\beta}\right) \cdot \left(\frac{\partial \theta_\gamma\left(\chi_\beta, \tau_\beta\right)}{\partial \chi} \boldsymbol{q_\gamma}\right) \tag{9b}$$
$$\boldsymbol{S_\beta} = \boldsymbol{S}\left(\boldsymbol{q_\beta}\right) \tag{9c}$$

where $\left(\chi_\beta, \tau_\beta\right)$ are the coordinates of the node corresponding to basis function θ_β.

If a modal basis is used, $\boldsymbol{F_\beta}, \boldsymbol{B_\beta}, \boldsymbol{S_\beta}$ may be found from the previous values of $\boldsymbol{q_\beta}$ in the iterative processes described below.

For functions $f(\chi,\tau) = f_\chi(\chi) f_\tau(\tau)$ and $g(\chi,\tau) = g_\chi(\chi) g_\tau(\tau)$, define the following integral operators:

$$[f,g]^t = f_\tau(t) g_\tau(t) \left\langle f_\chi, g_\chi \right\rangle \tag{10a}$$
$$\{f,g\} = \langle f_\tau, g_\tau \rangle \left\langle f_\chi, g_\chi \right\rangle \tag{10b}$$

Multiplying (5) by test function θ_α, using the polynomial approximations for $\boldsymbol{Q}, \boldsymbol{F}, \boldsymbol{B}, \boldsymbol{S}$, and integrating over space and time gives:

$$\left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \tau}\right\} \boldsymbol{q_\beta} = -\left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \chi}\right\} \boldsymbol{F_\beta} + \left\{\theta_\alpha, \theta_\beta\right\}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right) \tag{11}$$

2.1. The Discontinuous Galerkin Method

This method of computing the Galerkin predictor allows solutions to be discontinuous at temporal cell boundaries, and is also suitable for stiff source terms.

Integrating (11) by parts in time gives:

$$\left(\left[\theta_\alpha, \theta_\beta\right]^1 - \left\{\frac{\partial \theta_\alpha}{\partial \tau}, \theta_\beta\right\}\right) \boldsymbol{q_\beta} = [\theta_\alpha, \boldsymbol{w}]^0 - \left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \chi}\right\} \boldsymbol{F_\beta} + \left\{\theta_\alpha, \theta_\beta\right\}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right) \tag{12}$$

where $\boldsymbol{w}$ is the reconstruction obtained at the start of the time step with the WENO method. Define the following:

$$U_{\alpha\beta} = \left[\theta_\alpha, \theta_\beta\right]^1 - \left\{\frac{\partial \theta_\alpha}{\partial \tau}, \theta_\beta\right\} \tag{13a}$$
$$V_{\alpha\beta} = \left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \chi}\right\} \tag{13b}$$
$$\boldsymbol{W_\alpha} = \left[\theta_\alpha, \psi_\gamma\right]^0 \boldsymbol{w_\gamma} \tag{13c}$$
$$Z_{\alpha\beta} = \left\{\theta_\alpha, \theta_\beta\right\} \tag{13d}$$

Thus:

$$U_{\alpha\beta}\boldsymbol{q_\beta} = \boldsymbol{W_\alpha} - V_{\alpha\beta}\boldsymbol{F_\beta} + Z_{\alpha\beta}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right) \tag{14}$$

This nonlinear system in $\boldsymbol{q_\beta}$ is solved by a Newton method. The source terms must be solved implicitly if they are stiff. Note that $\boldsymbol{W}$ has no dependence on $\boldsymbol{q}$.

2.2. The Continuous Galerkin Method

This method of computing the Galerkin predictor is not suitable for stiff source terms, but it provides substantial savings on computational cost and ensures continuity across temporal cell boundaries.

$\{\psi_0, ..., \psi_N\}$ must be chosen in such a way that the first $N+1$ elements of $\left\{\theta_\beta\right\}$ have only a spatial dependence. The first $N+1$ elements of $\boldsymbol{q}$ are then fixed by demanding continuity at $\tau = 0$:

$$\boldsymbol{q}(\chi, 0) = \boldsymbol{w}(\chi) \tag{15}$$

where $\boldsymbol{w}$ is spatial the reconstruction obtained at the start of the time step with the WENO method.

For a given vector $\boldsymbol{v} \in \mathbb{R}^{(N+1)^2}$ and matrix $X \in M_{(N+1)^2,(N+1)^2}(\mathbb{R})$, let $\boldsymbol{v} = \left(\boldsymbol{v^0}, \boldsymbol{v^1}\right)$ and $X = \begin{pmatrix} X^{00} & X^{01} \\ X^{10} & X^{11} \end{pmatrix}$ where $\boldsymbol{v^0}, X^{00}$ are the components relating solely to the first $N+1$ components of $\boldsymbol{v}$. We only need to find the latter components of $\boldsymbol{q}$, and thus, from (11), we have:

$$\left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \tau}\right\}^{11} \boldsymbol{q_\beta^1} = \left\{\theta_\alpha, \theta_\beta\right\}^{11}\left(\boldsymbol{S_\beta^1} - \boldsymbol{B_\beta^1}\right) - \left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \chi}\right\}^{11} \boldsymbol{F_\beta^1} + \left\{\theta_\alpha, \theta_\beta\right\}^{10}\left(\boldsymbol{S_\beta^0} - \boldsymbol{B_\beta^0}\right) - \left\{\theta_\alpha, \frac{\partial \theta_\beta}{\partial \chi}\right\}^{10} \boldsymbol{F_\beta^0} \tag{16}$$

Define the following:

$$U_{\alpha\beta} = \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\tau}\right\}^{11} \tag{17a}$$

$$V_{\alpha\beta} = \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi}\right\}^{11} \tag{17b}$$

$$\boldsymbol{W_\alpha} = \left\{\theta_\alpha, \theta_\beta\right\}^{10}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right)^0 - \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi}\right\}^{10}\boldsymbol{F^0_\beta} \tag{17c}$$

$$Z_{\alpha\beta} = \left\{\theta_\alpha, \theta_\beta\right\}^{11} \tag{17d}$$

Thus:

$$U_{\alpha\beta}\boldsymbol{q^1_\beta} = \boldsymbol{W_\alpha} - V_{\alpha\beta}\boldsymbol{F^1_\beta} + Z_{\alpha\beta}\left(\boldsymbol{S^1_\beta} - \boldsymbol{B^1_\beta}\right) \tag{18}$$

Note that, as with the discontinuous Galerkin method, $\boldsymbol{W}$ has no dependence on the degrees of freedom in $\boldsymbol{q}$.

3. Conjecture

Extending the Galerkin method described in the previous section to three dimensions, the following system must be solved for $\boldsymbol{q}$:

$$\begin{aligned} U_{\alpha\beta}\boldsymbol{q_\beta} = \boldsymbol{W_\alpha} - V^1_{\alpha\beta}\boldsymbol{F_\beta} - V^2_{\alpha\beta}\boldsymbol{G_\beta} - V^3_{\alpha\beta}\boldsymbol{H_\beta} \\ + Z_{\alpha\beta}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right) \end{aligned} \tag{19}$$

where now we have the 3 scaled spatial variables χ_1, χ_2, χ_3 and $\boldsymbol{G}, \boldsymbol{H}$ are the flux components in the second and third spatial directions, respectively. In the case of the continuous Galerkin method, it is assumed that (19) is to be solved for only the non-fixed degrees of freedom in $\boldsymbol{q}$. The matrices $V^i_{\alpha\beta}$ are defined thus:

$$V^i_{\alpha\beta} = \left\langle\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi_i}\right\rangle \tag{20}$$

For the discontinuous Galerkin method, $\boldsymbol{W_\alpha}$ now takes the form:

$$\boldsymbol{W_\alpha} = \left[\theta_\alpha, \Psi_\gamma\right]^0 \boldsymbol{w_\gamma} \tag{21}$$

where $\Psi_\gamma(\chi_1, \chi_2, \chi_3)$ is an element of the following set, enumerated by γ:

$$\left\{\psi_i(\chi_1)\psi_j(\chi_2)\psi_c(\chi_3) : 0 \le i, j, k \le N\right\} \tag{22}$$

For the continuous Galerkin method, $\boldsymbol{W_\alpha}$ now takes the form:

$$\begin{aligned} \boldsymbol{W_\alpha} = \left\{\theta_\alpha, \theta_\beta\right\}^{10}\left(\boldsymbol{S_\beta} - \boldsymbol{B_\beta}\right)^0 \\ - \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi^1}\right\}^{10}\boldsymbol{F^0_\beta} - \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi^2}\right\}^{10}\boldsymbol{G^0_\beta} - \left\{\theta_\alpha, \frac{\partial\theta_\beta}{\partial\chi^2}\right\}^{10}\boldsymbol{H^0_\beta} \end{aligned} \tag{23}$$

Dumbser et al. [4] remark that for the continuous Galerkin case, the eigenvalues of $U^{-1}V^i$ are all 0 for $0 \le N \le 5$, for $i = 1, 2, 3$. Dumbser and Zanotti [7] state the same result for the discontinuous Galerkin case. This implies that in the conservative, homogeneous case ($\boldsymbol{B} = \boldsymbol{S} = \boldsymbol{0}$), owing to the Banach Fixed Point Theorem, existence and uniqueness of a solution are established, and convergence to this solution is guaranteed. As noted in Dumbser and Zanotti [7], in the linear case it is implied that the iterative procedure converges after at most $N + 1$ iterations.

In Dumbser et al. [4] it is conjectured that the result concerning the eigenvalues of $U^{-1}V^i$ holds for any N, and any number of spatial dimensions. A solution to this conjecture is provided here. For the theory in linear algebra required for this section, please consult a standard textbook on the subject, such as Nering [9].

3.1. The Discontinuous Galerkin Case

First, given the basis polynomials $\{\psi_0, \ldots, \psi_N\}$, define the following matrices:

$$\aleph_{ij} = \left\langle\psi_i, \psi_j\right\rangle \tag{24a}$$

$$\beth_{ij} = \left\langle\psi_i, \psi'_j\right\rangle \tag{24b}$$

Note that $\aleph$ is the Gram matrix, which by linear independence of $\{\psi_0, ..., \psi_N\}$ is invertible. Note also that if $p \in P_N$ then $p = a_j\psi_j$ for some unique coefficient vector $\boldsymbol{a}$. Thus, taking inner products with ψ_i, we have $\left\langle\psi_i, \psi_j\right\rangle a_j = \langle\psi_i, p\rangle$ for $i = 0, ..., N$. This produces the following result:

$$p = a_j\psi_j \Leftrightarrow \boldsymbol{a} = \aleph^{-1}\boldsymbol{x},\ x_i = \langle\psi_i, p\rangle \tag{25}$$

Without loss of generality, take the ordering:

$$\alpha = \alpha_t(N+1)^3 + \alpha_x(N+1)^2 + \alpha_y(N+1) + \alpha_z \tag{26}$$

where $0 \le \alpha_t, \alpha_x, \alpha_y, \alpha_z \le N$. Using the same ordering for β, we have:

$$U_{\alpha\beta} = \left(\psi_{\alpha_t}(1)\psi_{\beta_t}(1) - \beth_{\beta_t\alpha_t}\right)\cdot\aleph_{\alpha_x\beta_x}\cdot\aleph_{\alpha_y\beta_y}\cdot\aleph_{\alpha_z\beta_z} \tag{27a}$$

$$V^1_{\alpha\beta} = \aleph_{\alpha_t\beta_t}\cdot\beth_{\alpha_x\beta_x}\cdot\aleph_{\alpha_y\beta_y}\cdot\aleph_{\alpha_z\beta_z} \tag{27b}$$

Thus:

$$U = C \otimes \aleph \otimes \aleph \otimes \aleph \tag{28a}$$
$$V^1 = \aleph \otimes \beth \otimes \aleph \otimes \aleph \tag{28b}$$

where $C_{ij} = \psi_i(1)\psi_j(1) - \beth_{ji}$. Thus:

$$U^{-1}V^1 = \left(C^{-1}\aleph\right) \otimes \left(\aleph^{-1}\beth\right) \otimes I \otimes I \tag{29}$$

Therefore:

$$\left(U^{-1}V^1\right)^k = \left(C^{-1}\aleph\right)^k \otimes \left(\aleph^{-1}\beth\right)^k \otimes I \otimes I \tag{30}$$

A matrix X is nilpotent ($X^k = 0$ for some $k \in \mathbb{N}$) if and only if all its eigenvalues are 0. The conjecture will be proved if it is shown that $\left(\aleph^{-1}\beth\right)^k = 0$ for some $k \in \mathbb{N}$, as this would imply that $\left(U^{-1}V^1\right)^k = 0$, and thus all eigenvalues of $U^{-1}V^1$ are 0.

Taking $\boldsymbol{a} \in \mathbb{R}^{N+1}$, define:

$$p = a_0\psi_0 + \ldots + a_N\psi_N \in P_N \tag{31}$$

Note that:

$$(\beth \boldsymbol{a})_i = \left\langle \psi_i, \psi_0' \right\rangle a_0 + \ldots + \left\langle \psi_i, \psi_N' \right\rangle a_N = \left\langle \psi_i, p' \right\rangle \tag{32}$$

Thus, by (25):

$$\left(\aleph^{-1}\beth \boldsymbol{a}\right)_i \psi_i = p' \tag{33}$$

By induction:

$$\left(\left(\aleph^{-1}\beth\right)^k \boldsymbol{a}\right)_i \psi_i = p^{(k)} \tag{34}$$

for any $k \in \mathbb{N}$. As $p \in P_N$, $\left(\aleph^{-1}\beth\right)^{N+1} \boldsymbol{a} = \boldsymbol{0}$. As $\boldsymbol{a}$ was chosen arbitrarily, $\left(\aleph^{-1}\beth\right)^{N+1} = 0$. Thus, the conjecture is solved.

This proof is easily adapted to show that $U^{-1}V^2$ and $U^{-1}V^3$ are nilpotent, and clearly extends to any number of spatial dimensions. No specific choice has been made for $N \in \mathbb{N}$ and thus the result holds in general.

3.2. The Continuous Galerkin Case

In addition to $\aleph, \beth$, we now define $\aleph', \beth'$ where each new matrix is equal to the original, with its first row and column removed (the row and column corresponding to the constant-term polynomial ψ_0). Take the following ordering:

$$\alpha = \alpha_t (N+1)^3 + \alpha_x (N+1)^2 + \alpha_y (N+1) + \alpha_z \tag{35}$$

where now $0 \leq \alpha_x, \alpha_y, \alpha_z \leq N$ and $0 \leq \alpha_t \leq N-1$. Using the same ordering for β, we now have:

$$U_{\alpha\beta} = \beth'_{\alpha_t\beta_t} \cdot \aleph_{\alpha_x\beta_x} \cdot \aleph_{\alpha_y\beta_y} \cdot \aleph_{\alpha_z\beta_z} \tag{36a}$$
$$V^1_{\alpha\beta} = \aleph'_{\alpha_t\beta_t} \cdot \beth_{\alpha_x\beta_x} \cdot \aleph_{\alpha_y\beta_y} \cdot \aleph_{\alpha_z\beta_z} \tag{36b}$$

The proof for the continuous case follows in the same manner as the proof for the discontinuous case, with:

$$U^{-1}V^1 = \left((\beth')^{-1}\aleph'\right) \otimes \left(\aleph^{-1}\beth\right) \otimes I \otimes I \tag{37}$$

4. References

References

[1] Balsara, D. S., Meyer, C., Dumbser, M., Dub, H., Xu, Z., 2013. Efficient implementation of ader schemes for euler and magnetohydrodynamical flows on structured meshes - speed comparisons with runge-kutta methods. Journal of Computational Physics 235, 934–969.

[2] Balsara, D. S., Rumpf, T., Dumbser, M., Munz, C.-D., 2009. Efficient, high accuracy ader-weno schemes for hydrodynamics and divergence-free magnetohydrodynamics. Journal of Computational Physics 228, 2480–2516.

[3] Butcher, J. C., 2009. On fifth and sixth order explicit runge-kutta methods - order conditions and order barriers. Canadian Applied Mathematics Quarterly 17.

[4] Dumbser, M., Balsara, D. S., Toro, E. F., Munz, C. D., 2008. A unified framework for the construction of one-step finite volume and discontinuous galerkin schemes on unstructured meshes. Journal of Computational Physics 227 (18), 8209–8253.

[5] Dumbser, M., Castro, M., Pares, C., Toro, E. F., 2009. Ader schemes on unstructured meshes for nonconservative hyperbolic systems: Applications to geophysical flows. Computers & Fluids 38 (9), 1731–1748.

[6] Dumbser, M., Hidalgo, A., Zanotti, O., 2014. High order space-time adaptive ader-weno finite volume schemes for non-conservative hyperbolic systems. Computer Methods in Applied Mechanics & Engineering 268, 359–387.

[7] Dumbser, M., Zanotti, O., 2009. Very high order pnpm schemes on unstructured meshes for the resistive relativistic mhd equations. Journal of Computational Physics 228 (18), 6991–7006.

[8] Dumbser, M., Zanotti, O., Hidalgo, A., Balsara, D. S., 2013. Ader-weno finite volume schemes with space-time adaptive mesh refinement. Journal of Computational Physics 248, 257–286.

[9] Nering, E. D., 1970. Linear Algebra and Matrix Theory. John Wiley & Sons.

[10] Toro, E. F., 2009. Riemann Solvers and Numerical Methods for Fluid Dynamics: A Practical Introduction. Springer.

5. Acknowledgments

I acknowledge financial support from the EPSRC Centre for Doctoral Training in Computational Methods for Materials Science under grant EP/L015552/1.